

DOUBLE-TAKE® AVAILABILITY

Double-Take Availability provides affordable data protection, ensures minimal data loss and enables immediate recovery from any system outage. Double-Take Availability continuously captures changes as they happen and replicates those changes to one or more servers at any location, over any distance, so you always have access to a current copy of your data, applications and operating system. You can replicate to a disaster recovery site as far away as you would like, over standard IP networks, for maximum protection against data loss and improve performance by compressing the protected data before sending.

With Double-Take Availability, you can implement failover clusters without shared storage or geographic limitations – eliminating the single point of failure and providing the freedom to locate cluster nodes anywhere; a software approach encouraged by Microsoft®, Double-Take Availability ensures maximum availability for your Windows Server® clusters.

If you are consolidating workloads with Windows Server 2008 Hyper-V™ or VMware® vSphere, Double-Take Availability can replicate virtual machines from one virtualization host to another in real time and automatically fail over those virtual machines to the host in the event of an outage or disaster.

Versions of Double-Take Availability are available for Windows and Linux, support any application running on those OSs and work on physical or virtual servers. It is possible to replicate entire virtual machines from one host to another for VMware® vSphere and Microsoft® Hyper-V™ environments; regardless of the application or platform, you can ensure critical workloads are always available.

Intelligent Compression. Double-Take Availability is the only host-based replication solution that provides multilevel intelligent compression. Double-Take Availability provides four options for data compression that can be individually configured for each replication set defined within Double-Take Availability, allowing even further customization by server, data, and network.

Choose 'no compression' for less critical data or data that resides on application servers which require full availability of system resources or where network bandwidth is not a concern.

The next three levels enable compression, but to differing degrees. Higher compression levels require Double-Take Availability to use additional system resources (including CPU and RAM) on the source and target servers. Level one provides the least amount of compression, using the least amount of system resources and resulting in a smaller data transfer benefit. Level three performs the highest degree of compression, resulting in the least amount of data being transmitted; however, additional CPU and RAM is required. By providing multiple levels to choose from and allowing each to be set on individual replication sets, you can decide what level is necessary and what trade-offs are acceptable.

As data types can increase in size when compressed, resulting in more data being transmitted rather than less, Double-Take Availability compares the size of the compressed data and transmits the smaller of the two, ensuring that compression does not increase the load on the network. You don't have to be concerned with evaluating data for 'compressibility' before enabling compression – Double-Take Availability compresses only the data from which a benefit results.

Email Alerts. Double-Take Availability provides event notifications via email for immediate awareness of possible breaks in service levels. Email alerts can be configured with recipients for each server, each having its own event notification level (informational, warning, error). The email message includes useful information in the subject line, including server name where the event occurred, the error level and code. Administrators can easily manage events via email clients by sorting or filtering by these events, then choosing which ones require attention.

Asynchronous, Byte-level, Continuous Replication*. Asynchronous replication ensures that the replication process does not impact production applications. Double-Take Availability captures data for processing but does not keep it from updating the local disk. Double-Take Availability lets the application process data as it normally would and captures the changes for processing in the background. Double-Take Availability copies only the byte-level

*Except vSphere

changes to files, not entire blocks (physical or logical) or entire files, reducing the overhead on the servers and the network. Byte-level replication ensures that all transactions are captured and written in order on the target.

Hardware Independence. Hardware-based synchronous replication is often proprietary for each storage vendor and provides no cross-vendor replication, but Double-Take Availability has no affinity to any hardware or storage provider or storage architecture (SCSI, FibreChannel, iSCSI). This independence allows Double-Take Availability to be used on existing heterogeneous storage infrastructures and does not lock you in to a vendor for future purchases. Using a hardware-independent solution allows uses not possible with proprietary array-based solutions.

Streamlined Management Console. The simplicity of navigating Double-Take Availability lends itself to the proficient and extended use in your environment, reducing overall cost-of-ownership and improving ROI. Double-Take Availability provides an interface that is easy to use, ensuring you can easily protect your data. The Double-Take Availability UI automatically discovers servers running Double-Take Availability and displays them in a single window.

Uses Existing Networks. You do not need a private network for replicating data. Usually existing networks are more than sufficient, letting you implement lower-cost data replication and protection solutions, reducing the overall cost of ownership and removing any restrictions and additional costs for future network changes or purchases. Double-Take Availability can be run on its own private network should you want to isolate replication traffic.

Unlimited Distance Replication Double-Take. Availability uses standard IP networks so there are no distance limitations. The ability to replicate over long distances lends itself to numerous uses that short distance products cannot accomplish. Solutions such as centralized backup are possible and practical regardless of geographic separations. Double-Take Availability also offers protection against regional failures, replicating data across the country or across the globe to ensure the data is always available if needed.

Automated Failover. Double-Take Availability provides high-availability failover of servers to ensure users remain online in case of a failure. For customer-facing systems, it ensures that customer satisfaction remains high and sales are not impacted.

Open-file Mirroring and Replication. Applications do not need to be restarted when additional files and/or directories are configured for replication. Users, customers and applications remain online and active while Double-Take Availability is at work. Double-Take Availability can process open files and ensure they are fully replicated without taking the files offline.

Bandwidth Throttling. User-defined controls within Double-Take Availability can limit the amount of available network bandwidth it can use for data replication. This allows replication to occur real-time without affecting users on the same network. And, since Double-Take Availability can queue data for transmission, all changes are updated to the target and not lost due to network limits. Double-Take Availability lets you define the actual type of network connection being used (T1, 128Kbps, etc.) and then define the amount that can be used for replication. Double-Take Availability will not go beyond that limit, regardless of the amount of data it has to transfer. Double-Take Availability lets you limit network use during busy work hours and increase or remove limits during non-peak hours.

Failback/Restore. Should a failure occur, Double-Take Availability can restore data from the target back to the original source or to an alternate location. Through the Double-Take Availability UI, you can restore data from the replicated disk back to the production disk once the failure is corrected, greatly reducing the time to recover and restore and ensuring that you recover from the time the failure occurred, not from when the last backup was taken. Unlike other solutions that make you remember which files came from what location, the Double-Take Availability restore process automatically reverses the direction of your original replication job.

Block Checksum Re-Mirror. Should a disconnect occur between the source and target, instead of doing a complete mirror of the entire replication set, Double-Take Availability can perform a block-checksum re-mirror. This re-mirror replicates the file differences between the source and target, which takes much less time and resources to accomplish, ensuring that the target is coordinated with the source.

Replication Scheduling. Double-Take Availability provides the flexibility to schedule when replication occurs. Double-Take Availability continues to filter and queue up all appropriate data changes until the

scheduled time has arrived, then replicates the queued changes to the target until the user-defined replication window closes. This lets you fine tune Double-Take Availability around business needs to ensure network and system resources are used efficiently and that replication does not impact production.

Reports and Statistics. Double-Take Availability provides reports and statistics about your replication environment that let you know what is happening on a Double-Take Availability server, all from a single interface. Reports include information such as what was last replicated, how much has been replicated, failovers, downed systems, and other useful statistics for ensuring adequate data protection.

64-bit Support. Double-Take Availability provides 64-bit versions that run on Windows® Server x64 Edition (AMD64 and EMT64) or Windows Server 2003 for Itanium-based systems.

One-to-Many Replication. To ensure all nodes in a cluster can allow failover from any other member node it is crucial that the data is available on all nodes. Double-Take Availability lets each cluster member replicate data to any or all members of the cluster, ensuring that the data is available and current in an outage.

Many-to-One Failover. Double-Take Availability can be configured so that a single target server is used for numerous source servers, reducing the total cost of ownership. Single target configuration also helps facilitate centralized backup since many local or remote servers can be backed up from a single server, further lowering costs by reducing the number of backup licenses needed.

Selectable Replication Routes. On target servers with multiple network adapters, you can specify the IP address of the network adapter that Double-Take Availability uses to replicate data to the target.

Flow Control*. Double-Take Availability is designed to handle spikes in the data rate of change even if the network connection to the target server is not sufficient to handle all data simultaneously. Double-Take Availability continues to filter all file changes and buffer these changes, while simultaneously transmitting to the target as quickly as possible. Queuing ensures that all transactions are replicated to the target without data loss. With bandwidth throttling

and scheduling features, users can configure Double-Take Availability around production server and network requirements. If performing backups from the target sever, replication can be paused to ensure a complete point-in-time backup, while ensuring that all changes on the source server are still transmitted to the target and are applied once the backup is complete.

Automatic Re-mirror*. Should a problem or scheduled maintenance require that the connection between the source and target be broken, Double-Take Availability will re-establish its connection when possible and will automatically re-mirror the source to the target, ensuring it remains in sync with the source.

Full Command-Line Control*. Double-Take Availability allows all GUI functions to be controlled from the commandline, via scripts or individual commands, giving you greater flexibility and automation capabilities.

Target Reboot without Re-mirror*. You can gracefully shutdown the target server without re-synchronizing the entire protected dataset.

Windows-Specific Features

Livewire Option. Using the same technology as full-server failover, the Livewire option provides many-to-one full server protection, with the ability to recover servers as needed after failure. The Livewire option supports 32bit or 64-bit servers and can recover over the LAN or WAN to dissimilar hardware.

Virtual Recovery Assistant. The Virtual Recovery Assistant lets you protect a source physical or virtual machine to a virtual target machine running on either Microsoft Windows Server 2008 with Hyper-V or VMware ESX Server. Data is replicated directly into the virtual machine's virtual disk, which means that the target virtual machine can remain powered off until it is needed for a failover, allowing over-subscription of the virtual environment at the recovery site and reducing licensing fees for guest operating systems and applications. The Virtual Recovery Assistant provisions the target virtual machine for you and monitors the protected source machine. During protection setup, Virtual Recovery Assistant also lets you increase the drive size, memory, and number of processors on the target virtual

*Except vSphere

machine. It can be configured to failover to the target automatically if the connection is lost between the source and target.

Storage Optimization. The Cargo feature optimizes storage on file servers running Double-Take Availability. Cargo lets you improve mirror performance, reduce disk space used and select restore speeds. Cargo reduces any size file to 4K and creates a pointer to a full-sized file on the target server. Cargo provides the ability to setup regularly scheduled scans of repsets. Additionally, you can begin an immediate archive of files based on the age and size of a file and select individual folders or files from a Double-Take Availability repset for immediate archive or recall.

Enterprise Install Console. The Enterprise Install Console helps you deploy and maintain Double-Take Availability installations. From within this console, you can automate installs or upgrades on all servers throughout your enterprise.

Single-Screen Server Monitoring with Dashboard. The Double-Take Availability Dashboard provides the ability to monitor the health of Double-Take Availability source servers on one screen. The screen displays red, yellow or green visual indicators of status and is updated continually based on the monitoring interval selected by the user. Double-Take Availability servers present in the management console are imported into the Dashboard to let you quickly and easily begin monitoring your protected servers.

The Dashboard provides one-click sort and filter for monitored source servers. You can sort and filter by server name, repset name, mirror status, repset status, disk queue and many other key variables. Select a monitored server and open other features such as the Management Console or Application Manager with one click. Additionally, the Dashboard provides the ability to view event log information and mirror status.

Reporting Center. The Reporting Center is a centralized reporting and analysis feature that offers detailed, custom analysis and reports of your entire Double-Take Availability environment. Reporting Center collects and analyzes protection statistics and provides rich reports and dashboard views to show the overall effectiveness of your protection scheme to key stakeholders in your organization.

The Reporting Center lets you show the overall effectiveness of your protection scheme over time, giving you the information that you need to be accountable to management. It also lets you verify the integrity of the data on your target server and identify subtle or growing problems that may only be visible when looking at trends.

Active/Active MSCS Support. If running in a Microsoft clustered environment in an Active/Active configuration, Double-Take Availability can be configured to replicate data from any or all members, to each other or to remote systems. Double-Take Availability can be configured to be aware of failovers and automatically adjust the sources and targets accordingly, removing the need for user intervention at the time of failure.

Full-Server Failover. The Full-Server Failover feature applies the source server's OS configuration, applications and data to the target server. Because applications need not be pre-installed and maintained on the target server, full-server failover is easy to configure. Full-server failover supports one-to-one connections of 32bit or 64-bit servers and can failover the LAN or WAN to dissimilar hardware.

Windows Server Failover Clusters without Shared Storage. Enhance the availability of clustered applications by removing the shared disk single-point-of-failure. Double-Take Availability works with Windows Server Failover Clustering to provide highly available applications by allowing the configuration of separate storage on each of the cluster members. Windows Server Failover Clustering requires clustered disk resources to be on a shared disk subsystem; should the disk storage fail for any reason, the entire cluster becomes unavailable. Maintaining separate storage removes the risk of cluster or application failure due to a single disk failure. Double-Take Availability ensures that the data on the active member node and disk are replicated to the member nodes in a failover.

Long Distance Clusters. With SCSI distance limitations, cluster members have distance limitations (they all must connect to the same shared-storage devices). Because Double-Take Availability removes the shared storage requirement, cluster members can be separated over long distances to provide disaster recovery and high availability in a single solution. In a failure, the applications will failover to an active member, and since that member has its own copy of the data, the application is immediately available.

Application Support (GeoCluster feature).

Double-Take Availability provides support for all Windows Server cluster-aware applications, including Hyper-V. Double-Take Availability uses the Windows Server Failover Clustering API to provide native support for cluster-aware applications to ensure full functionality and compatibility.

Online Pending Logic (GeoCluster feature).

Double-Take Availability protects data from corruption by stopping a cluster from failing over to a node that is suspect. If the target is suspect, the Online Pending state provides the option to verify data on the target before putting it online, revert to the last known good state on the target, or force the resource offline.

File Share Witness Support (GeoCluster feature).

Microsoft Cluster Service's File Share Witness is an option for quorum management. (File Share Witness is only available for 2-node clusters.)

Majority Node Set (GeoCluster feature).

Windows Server Failover Clustering needs to maintain quorum devices in order to know when a cluster member owns a resource. Because Double-Take Availability splits up the storage resources unbeknownst to Windows Server, Double-Take Availability has to perform additional processing and replication in order to provide the quorum resources needed. Windows Server can be configured to maintain quorum on its own even when using non-shared storage devices, allowing the clustering application to handle this on its own.

Integrated User Interface (GeoCluster feature).

Double-Take Availability replication set configuration is integrated in the Windows Server Failover Clustering Administrator, further simplifying ease of management. The Double-Take management console can be used for advanced monitoring of connections and events within Double-Take Availability.

Windows and Linux-Specific Features

Management Console Server Filtering. Server filtering lets you create custom server views within the Double-Take Availability management console. By hiding servers, you can see a select set of servers, simplifying manageability of the environment. Views are created by individual login IDs, allowing each administrator to see only the servers with which they are concerned. Using customized server

views, Double-Take Availability allows for virtually unlimited scalability, as you can choose to view as few servers as desired without having to search through other servers.

Product Update Checking. The installation process lets you check (via the web) for the latest product updates and releases. You are given a list of updates available and can choose to download and install immediately. This update check is available via the Double-Take Availability GUI, allowing you to check for updates anytime. Double-Take Availability determines what updates are installed and verifies that these are all included before presenting it as an option for download. Any update that does not include all patches installed by the customer is not listed.

Task Command Processing. Double-Take Availability lets you insert commands into the data stream for execution at points during regular processing. Tasks such as automatically initiating target server backup are possible. Via in-band commands, you can ensure all files on the target are in sync with each other, pause writing on the target and initiate the backup. Once the backup is complete the task command can enable writing on the target.

Resource Planning Tool. To ensure an optimal replication environment, Double-Take Availability includes a Resource Planning Tool that runs a filter process without actually copying the data to a remote target, providing information necessary to implement the networks and servers to meet your RPO and RTO. By running this tool before deployment, you can avoid overspending on hardware and networks and avoid an insufficient deployment.

Serverless Backup Support. By replicating production data to a target server, backups can be performed from the target server without having to shut down applications or requiring users to log off the production server. The backup window is no longer an issue, as the target server has virtually no time limits for the backup to complete. Users continue to run on the production server while backup is offloaded to non-production target servers.

File/Directory Selection. You can select full file systems, directories or files for replication, giving you greater flexibility to configure efficient use of resources. Important files can be selected, while system and temporary files can be skipped. The use of wildcards and drag-n-drop facilitate quick and simple configuration.

Support for Different File Types. Double-Take Availability has been successfully tested against file types such as encrypted, compressed and ultra-long file names.

Verification. Verification can be run to ensure the target is in sync with the source, which is useful if there is a brief outage, if services are stopped, or if someone has made direct updates to the data on the target.

SNMP Counters and Traps. Double-Take Availability simplifies management by integrating with your enterprise management frameworks and forwarding replication statistics and events via SNMP.

Enhanced Restores. Restores in Double-Take Availability maximize uptime for users because there is no need to take applications offline during a restore. Any changes made by users between the time the restore is started and initiation of failback are replicated to the source.

Advanced Scripting. Script execution points activate scripts around mirror events such as mirror start, mirror complete and mirror stop.

Server Groups. Server groups and health status information within the Double-Take Availability console let you manage multiple Double-Take Availability servers and centrally monitor the health of real-time replication and failover.

ICMP-less Ping. Additional options for failover monitoring provide the ability to monitor the Double-Take Availability service in place of monitoring the IP address of the production server.

MOM Integration Pack. Double-Take Management Pack for Microsoft Operations Manager lets you monitor issues on all Double-Take Availability servers across your enterprise and alerts you to only important events, letting you stay aware of, respond to, and demonstrate accountability for server and application services levels.

Volume Shadow Copy Service Snapshot Integration. Double-Take Availability integrates with Microsoft Volume Shadow Copy Service to let you schedule and recover from up to 64 point-in-time copies of data on your Double-Take Availability target, ensuring that you can recover from human error, corruption, viruses and other unwanted changes.

Interoperability with HSM Products. Double-Take Availability interoperates with Hierarchical Storage Management solutions from Commvault, Bridgehead and DiskXtender. When these products archive files, replace them with placeholder files and recall the files, Double-Take Availability replicates their every move. Additionally, when Double-Take Availability performs Verify operations on archived files, it recognizes the placeholder files and avoids unwanted file recalls.

Target State Awareness. Double-Take Availability can report the "state" of replicated data, providing additional status information about the replication process.

The VSS Snapshot Integration feature uses target state awareness to take automatic point-in-time snapshots of protected data as data moves from a "known" to an "unknown" state, ensuring that the administrator has at least a "known-good" copy of protected data from which to recover.

Target Path Blocking. Block any other application or process from writing to the replicated data on the Double-Take Availability target server, preventing unwanted changes and ensuring that replicated data does not become corrupt.

VMware-Specific Features

Agent-less Installation. For vSphere servers, no "agent" software is required, providing easy management and rapid deployment of protection scenarios. Install Double-Take Availability once and multiple protection scenarios can be launched. This approach aggregates resource consumption by VM protection so it can be monitored and managed at the VMware enterprise level.

Scheduled Start. Optimize network resources by configuring Double-Take Availability to start your initial vSphere synchronization at a future point in time. By default, Double-Take Availability starts the initial synchronization of a protection job immediately after the protection is created. To specify a date and time in the future at which the initial synchronization is started, a date and start time can be specified in the protection wizard.

One-Click Failover and Restore. Double-Take Availability lets you failover or restore a vSphere virtual machine with one click. When selected from the monitor screen, the "Failover" function shuts

down the source virtual machine, starts the replica virtual machines and begins the failover process. When a protection job that is in a failed state is

selected, a job is configured to synchronize the virtual machine back to the production vSphere host. When the restoration completes, the administrator can "failback" automatically. There are three options for failover:

- *Live failover* – perform a failover with network connectivity
- *Test failover* – perform a test failover in which the source is not shut down in order to verify data integrity on the target
- *Undo failover* – lets you re-start the source virtual machine in the state it was in at the point of failover and discard any data changes that may have taken place on the target vSphere host.

File Difference Mirror. Double-Take Availability supports file difference mirrors at the block level during a re-mirror of the vSphere virtual machine. Only the changed blocks of the file are sent to the target, increasing mirroring speed and decreasing bandwidth requirements.

Support for Multiple VirtualCenter Servers. Double-Take Availability lets you protect vSphere virtual machines when the source and target vSphere hosts are managed by different VirtualCenter servers.

Standalone Installation Support. Double-Take Availability can communicate with vSphere hosts either directly or through VMware VirtualCenter. This allows for the deployment of Double-Take Availability in environments where VirtualCenter is not installed.

Target Path Mapping. Double-Take Availability lets you specify the path used to create the target vSphere virtual machine.

Pre-Staged Target Data Support. Double-Take Availability keeps track of all the information about the association of any vSphere source .vmdk file to the corresponding target .vmdk file. The mapping information can be used to perform target pre-staging. Pre-staging means before creating a job, you would copy information through other means, and after a job is created, Double-Take Availability uses a difference mirror instead of a full mirror.

Setup Wizard. Setting up a vSphere protection scenario is easy with a directed step-by-step workflow

that guides you through the process. With a few mouse clicks, and in less than a minute, a VM protection scenario can be created and deployed.

Virtual Infrastructure Browsing. The Double-Take Availability management interface uses the logical groupings and hierarchies available in VMware VirtualCenter, making it possible to navigate through configured virtual machines and select them for management and protection.

Data Compression Options. Double-Take Availability provides options for compressing the vSphere replicated files, providing flexible trade offs between host system resource consumption and network bandwidth usage. Available compression options include:

- *Never* – No compression is applied to the data.
- *Only during the initial synchronization* – Data is compressed only during the initial transmission of the flat virtual disk files.
- *Always* – All transmissions associated with the protection set is compressed.

Smart Resume. In case of an interruption on during the initial synchronization, the vSphere virtual disk version is "remembered" (via a date / time stamp) so when protection is resumed, no re-synchronization is needed (providing the target is the same version as it was at the start). Once protection is resumed, you can move directly to monitoring and replication.

Replication Frequency Options. For vSphere servers, Double-Take Availability provides flexibility to configure the frequency of change replication based on the amount of change or the time that has elapsed since the last replication window.

Hyper-V Specific Features

Auto Discovery. Built to use the Microsoft Hyper-V management interfaces (APIs), Double-Take Availability automates the process of finding and protecting all files associated with Microsoft Hyper-V virtual machines, including configuration files, virtual hard disks (VHDs) and snapshots. Double-Take Availability lets you see a listing of all virtual machines in their environment and to select which workloads to protect.

Auto Provisioning. Double-Take Availability provides failover for selected virtual machines. When a virtual machine is failed over, Double-Take

Availability automatically provisions the virtual machine on a new server, and runs the virtual machine with all pre-existing snapshots configured from the original virtual machine.

Virtual Switch Mapping. Map your source virtual machine NICs to a virtual switch on the target Hyper-V host. At failover, your server connects to the target virtual switch.

Unlimited Protection. Double-Take Availability can protect a single virtual machine, or all virtual machines in an environment, configure replication jobs per virtual machine and monitor and failover all machines independently of one another. Two licenses of Double-Take Availability, (one for the source and

one for the target) can be used to protect all virtual machines on both machines.

Test Failover with Undo. Double-Take Availability lets you stop replication and power on the replica virtual machine with no network connection, allowing the replica virtual machine be tested while the original source VM is still running.

Guest OS Agnostic. Double-Take Availability can be used to protect virtual machines running any operating system supported by Microsoft Hyper-V.

Target Path Mapping. Double-Take Availability lets you specify the path used to create the target Hyper-V virtual machine.

Manage your subscription to eNews. Visit: www.doubletake.com/subscribe

Double-Take Availability™ 

 Printed on recycled paper.

Get the standard today: www.doubletake.com or 888-674-9495

© Double-Take Software, Inc. All rights reserved. Double-Take, Balance Double-Take Cargo, Double-Take Flex, Double-Take for Hyper-V, Double-Take for Linux, Double-Take Move, Double-Take ShadowCaster, Double-Take for Virtual Systems, GeoCluster, Livewire, netBoot/i, NSI, sanFly, TimeData, TimeSpring, winBoot/i and associated logos are registered trademarks or trademarks of Double-Take Software, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries. Microsoft, Hyper-V, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. Novell, the Novell logo, the N logo, SUSE are registered trademarks of Novell, Inc. in the United States and other countries. All other trademarks are the property of their respective companies.